arXiv:cs/0410039v1 [cs.DS] 17 Oct 2004

Generating All Maximal Induced Subgraphs for Hereditary, Connected-Hereditary and Rooted-Hereditary Properties

Sara Cohen* Yehoshua Sagiv*

Abstract

The problem of computing all maximal induced subgraphs of a graph G that have a graph property $\mathcal{P}$, also called the *maximal $\mathcal{P}$-subgraphs problem*, is considered. This problem is studied for hereditary, connected-hereditary and rooted-hereditary graph properties. The maximal $\mathcal{P}$-subgraphs problem is reduced to restricted versions of this problem by providing algorithms that solve the general problem, assuming that an algorithm for a restricted version is given. The complexity of the algorithms are analyzed in terms of total polynomial time, incremental polynomial time and the complexity class P-enumerable. The general results presented allow simple proofs that the maximal $\mathcal{P}$-subgraphs problem can be solved efficiently (in terms of the input and output) for many different properties.

1 Introduction

Hereditary and connected-hereditary graph properties include many common types of graphs such as cliques, bipartite graphs and trees. Such properties appear in many contexts, and thus, they have been widely studied, e.g., [14, 20, 3, 17]. This paper focuses on the *maximal $\mathcal{P}$-subgraphs problem*: Given a graph property $\mathcal{P}$ and an arbitrary graph G, find all maximal induced subgraphs of G that have the property $\mathcal{P}$. We consider properties that are hereditary, connected hereditary or rooted hereditary (a variant of connected hereditary).

Since the output for the maximal $\mathcal{P}$-subgraphs problem may be large, our complexity analysis takes into consideration both the size of the input and the size of the output. Specifically, we consider the complexity measures *total polynomial time* [10] and *incremental polynomial time* [10], and the complexity class *P-enumerable* [19].

The maximal $\mathcal{P}$-subgraphs problem has been studied for many properties $\mathcal{P}$. For example, it has been shown that this problem is both P-enumerable and solvable in incremental polynomial time for the properties "is an independent set" and "is a clique" [10, 18, 1]. If subgraphs (and not only induced subgraphs) are allowed, then the maximal $\mathcal{P}$-subgraphs problem is P-enumerable for the properties "is a spanning tree" [16, 13], "is an elementary cycle" [16, 9] and "is an elementary path" [16, 2], among others.

This paper differs from previous work in that we do not consider specific properties, but instead, deal with the problem for a general $\mathcal{P}$. Our strategy is to reduce the maximal $\mathcal{P}$-subgraphs problem to restricted versions of the maximal $\mathcal{P}$-subgraphs problem that are often easier to solve. Our

*The Selim and Rachel Benin School of Computer Science and Engineering, The Hebrew University of Jerusalem, Givat Ram Campus, Jerusalem 91904, Israel. Email: {sarina,sagiv}@cs.huji.ac.il

reductions are by means of algorithms that solve the maximal $\mathcal{P}$-subgraphs problem, given a solution to a restricted version of this problem.

Using our algorithms, we can show that the maximal $\mathcal{P}$-subgraphs problem can be solved in total-polynomial time, incremental-polynomial time or is P-enumerable if certain conditions hold on the runtime of a restricted version of this problem. Hence, our approach is easily shown to include and improve upon previous results in graph theory. For example, it is shown that the maximal "bipartite"-subgraph problem is solvable in total polynomial time, which improves upon [15].

The maximal $\mathcal{P}$-subgraphs problem also has immediate practical applications in the database field. Interestingly, it turns out that many well-known semantics for answering queries in the presence of incomplete information can be modeled as hereditary, connected-hereditary or rooted-hereditary graph properties. Hence, the results in this paper imply and improve upon the complexity results in [11, 12, 4] and imply the complexity result in [21]. In fact, modeling semantics as graph properties allows previously presented semantics to be extended without affecting their complexity. See [5] for more details.

2 Graphs and Graph Properties

Graphs and Induced Subgraphs. A *graph* $G = (V, E, r)$ consists of *(1)* a finite set of *vertices* V, *(2)* a set of *edges* $E \subseteq V \times V$ *and (3)* a *root* r such that $r \in V \cup \{\perp\}$. We say that G is *rooted* if *(1)* $r \neq \perp$ *and (2)* every vertex in G is reachable via a directed path from r. We say that G is *connected* if its underlying undirected graph is connected. Observe that every rooted graph is connected. However, a connected graph need not be rooted. We use $V(G)$ to denote the set of vertices of G.

A graph H is an *induced subgraph* of a graph G, written $H \subseteq_{\text{IS}} G$, if *(1)* H is derived from G by deleting some of the vertices of G (and the edges incident on these vertices) *and (2)* H has the same root as G, if the root of G is among the vertices of H, and has $\perp$ as its root otherwise. We write $H \subset_{\text{IS}} G$ if $H \subseteq_{\text{IS}} G$ and H is not equal to G.

We use $G[\{v_1, \ldots, v_n\}]$ to denote the induced subgraph of G that contains exactly the vertices $v_1, \ldots, v_n$. If H and H' are induced subgraphs of G and v is a vertex in G, we use $G[H]$, $G[H, v]$ and $G[H, H']$ as shorthand notations for $G[V(H)]$, $G[V(H) \cup \{v\}]$ and $G[V(H) \cup V(H')]$, respectively.

Graph Properties. A *graph property* $\mathcal{P}$ is a nonempty and possibly infinite set of graphs. For example, "is a clique" is a graph property that contains all graphs that are cliques. In this paper, we only consider properties $\mathcal{P}$ such that it is possible to verify whether a graph G is in $\mathcal{P}$ in polynomial time. Hence, we assume that there is a polynomial procedure $\textsc{Sat}\langle\mathcal{P}\rangle$ that receives a graph G as input, and returns true if $G \in \mathcal{P}$ and false otherwise. Observe that the notation $\langle\mathcal{P}\rangle$ denotes an algorithm that is parameterized by the graph property $\mathcal{P}$, i.e., that differs for each value of $\mathcal{P}$.

We consider several special types of graph properties. A graph property $\mathcal{P}$ is *hereditary* if $\mathcal{P}$ is *closed* with respect to induced subgraphs, i.e., whenever $G \in \mathcal{P}$, every induced subgraph of G is also in $\mathcal{P}$. A graph property $\mathcal{P}$ is *connected hereditary* if *(1)* all the graphs in $\mathcal{P}$ are connected *and (2)* $\mathcal{P}$ is closed with respect to connected induced subgraphs. A graph property $\mathcal{P}$ is *rooted-hereditary* if *(1)* $\mathcal{P}$ only holds on rooted graphs *and (2)* $\mathcal{P}$ is closed with respect to rooted induced subgraphs.[1] It is rather unusual for a rooted-hereditary property to also be hereditary

[1] This type of properties is useful when considering database problems related to semistructured data, since semistructured data are usually represented as rooted graphs.

or connected-hereditary. Actually, one can show that if $\mathcal{P}$ is rooted-hereditary and $\mathcal{P}$ is also hereditary or connected-hereditary, then $\mathcal{P}$ contains only graphs with at most one vertex. This gives an additional motivation to considering rooted-hereditary properties, since they generally differ from connected-hereditary and hereditary properties.

Many graph properties are hereditary [8], e.g., "is a clique" and "is a forest." Note that "is a clique" is also connected hereditary. However, "is a clique" is not rooted hereditary, since it contains graphs that do not have roots. Some properties are connected hereditary, but not hereditary or rooted hereditary, such as "is a tree," which contains a graph G if the underlying undirected graph of G is a tree. Note that G is not necessarily rooted. Hence, "is a tree" is not rooted hereditary. The property "is a rooted clique" is rooted hereditary.

The Maximal $\mathcal{P}$-Subgraphs Problem. Let G be a graph and $\mathcal{P}$ be a property. (The graph G is not necessarily in $\mathcal{P}$.) We say that H is a *$\mathcal{P}$-subgraph* of G if $H \subseteq_{\text{IS}} G$ and $H \in \mathcal{P}$. The set of $\mathcal{P}$-subgraphs of a graph G is denoted $\mathcal{P}(G)$.

We say that H is a *maximal $\mathcal{P}$-subgraph* of G if H is a $\mathcal{P}$-subgraph of G and there is no $\mathcal{P}$-subgraph H' of G, such that $H \subset_{\text{IS}} H'$. We use $\mathcal{P}^m(G)$ to denote the set of maximal $\mathcal{P}$-subgraphs of G. The *maximal $\mathcal{P}$-subgraphs problem* is: Given a graph G, find the set $\mathcal{P}^m(G)$.

3 Complexity Classes and Measures

This paper explores the problem of computing $\mathcal{P}^m(G)$, for a hereditary, connected-hereditary or rooted-hereditary property $\mathcal{P}$ and an arbitrary graph G. The maximal $\mathcal{P}$-subgraphs problem cannot be solved in polynomial time, in the general case. This follows from the fact that sometimes the size of $\mathcal{P}^m(G)$ is exponential in the size of G (see [3] for details). Hence, exponential time may be needed just to print the output. In this section, we discuss two complexity measures that are of interest when the output of a problem may be large: *total polynomial time* [10] and *incremental polynomial time* [10]. We also consider the complexity class *P-enumerable* [19].

A problem can be solved in *total polynomial time*, or PIO for short, if the time required to list all its solutions is bounded by a polynomial in n (the size of the input) and K (the number of solutions in the output).[2] For the maximal $\mathcal{P}$-subgraphs problem, n is the number of vertices in G and K is the number of graphs in $\mathcal{P}^m(G)$.

The complexity class *P-enumerable* is more restrictive than the measure of total polynomial time. Formally, a problem is P-enumerable if the time required to list all its solutions is bounded by K times a polynomial in n. Note that P-enumerable differs from total polynomial time in that the factor of the output in the runtime must be linear. Since the size of the output may be exponential in the size of the input, the factor of output size in the total runtime is highly influential.

Another complexity measure that is of interest when dealing with problems that may have large output (such as the maximal $\mathcal{P}$-subgraphs problem) is *incremental polynomial time*. Formally, a problem is solvable in incremental polynomial time, or PINC for short, if, for all k, the k-th solution of the output can be returned in polynomial time in n (the input) and k. Incremental polynomial time is of importance when the user would like to optimize evaluation time for retrieval of the first k maximal induced subgraphs, as opposed to optimizing for overall time. This is particularly useful

[2] This complexity measure is similar to polynomial time input-output complexity, which is commonly considered in database theory, e.g. [21].

in a scenario where the user reads the answers as they are delivered, or is only interested in looking at a small portion of the total result. If a problem is solvable in total polynomial time, but not in incremental polynomial time, the user may have to wait exponential time until the entire output is created, before viewing a single maximal $\mathcal{P}$-subgraph.

Observe that every problem that is P-enumerable is also solvable in total polynomial time. Similarly, every problem that is solvable in incremental polynomial time is also solvable in total polynomial time. It is not known whether every problem that is solvable in incremental polynomial time is also P-enumerable, and vice-versa. The maximal $\mathcal{P}$-subgraphs problem has been studied for many properties $\mathcal{P}$. See Section 1 for several examples and see [6] for a listing of algorithms for combinatorial enumeration problems.

4 Restricting the Maximal $\mathcal{P}$-Subgraphs Problem

Let $\mathcal{P}$ be a graph property. Suppose that we want to show that the maximal $\mathcal{P}$-subgraphs problem is in PIO. To do this we must devise an algorithm that, when given any graph G, produces $\mathcal{P}^m(G)$ in polynomial time in the input (i.e., G) and the output (i.e., $\mathcal{P}^m(G)$). For many properties $\mathcal{P}$, it is difficult to find such an algorithm, since an arbitrary graph G must be dealt with. Our task of finding an appropriate algorithm is even more difficult if we actually want to show that the maximal $\mathcal{P}$-subgraphs problem is P-enumerable or is in PINC. Hence, we focus on restricted versions of the maximal $\mathcal{P}$-subgraphs problem. For these restricted versions, it is often easier to devise an efficient algorithm. Later on we will show how, given an algorithm for one of the restricted problems, the general problem can be solved.

Let G be a graph and let $\mathcal{P}$ be a property. We use $G - v$ to denote the induced graph of G that contains all vertices other than v. We say that G *almost satisfies* $\mathcal{P}$ if there is a vertex v in G, such that $G - v \in \mathcal{P}$. Let v' be a vertex in G. We use $\mathcal{P}^m(G;\, v')$ to denote the subset of $\mathcal{P}^m(G)$ that contains graphs with the vertex v'.

We will be interested in three restricted versions of the maximal $\mathcal{P}$-subgraphs problem.

- The *input-restricted maximal $\mathcal{P}$-subgraphs problem* is: Given a graph G that almost satisfies $\mathcal{P}$, find all maximal $\mathcal{P}$-subgraphs of G.
- The *output-restricted maximal $\mathcal{P}$-subgraphs problem* is: Given an arbitrary graph G and a vertex v' in G, find all maximal $\mathcal{P}$-subgraphs of G that contain v'.
- The *io-restricted maximal $\mathcal{P}$-subgraphs problem* is: Given a graph G that almost satisfies $\mathcal{P}$ and given a vertex v' in G, find all maximal $\mathcal{P}$-subgraphs of G that contain v'.

Note that the output-restricted maximal $\mathcal{P}$-subgraphs problem can be used in a straightforward way to solve the maximal $\mathcal{P}$-subgraphs problem. However, it is not clear how the other two problems can be used to solve the maximal $\mathcal{P}$-subgraphs problem.

The complexity of these three problems is highly dependent on the graph property $\mathcal{P}$. Sometimes, it turns out that the input-restricted maximal $\mathcal{P}$-subgraphs problem and the io-restricted maximal $\mathcal{P}$-subgraphs problem can actually be solved in polynomial time, since the number of graphs in their output is bounded in size by a constant. However, the fact that G almost satisfies $\mathcal{P}$ does not always entail that $\mathcal{P}^m(G)$ is small. This is shown in the following example.

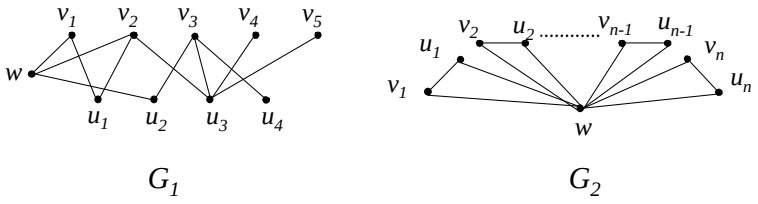

Figure 1: Finding Maximal Bipartite Subgraphs

Example 4.1. Let $\mathcal{P}_{\text{CBIP}}$ be the connected-hereditary property that contains all connected bipartite graphs. Suppose that G is a graph that almost satisfies $\mathcal{P}_{\text{CBIP}}$. It is not difficult to see that $\mathcal{P}^m_{\text{CBIP}}(G)$ contains at most three graphs. As an example, consider the graph G_1 in Figure 1. The graph G_1 almost satisfies $\mathcal{P}_{\text{CBIP}}$, since $G_1 - w$ is a connected bipartite graph. The set $\mathcal{P}^m_{\text{CBIP}}(G_1)$ contains the following three graphs: *(1)* $G_1 - w$, *(2)* $G_1 - u_2$ (derived by removing the neighbors of w on the bottom side) *and (3)* $G_1 - \{v_1, v_2, u_1\}$ (derived by removing the neighbors of w on the top side, and then removing unconnected vertices).

Let $\mathcal{P}_{\text{BIP}}$ be the hereditary property that contains all bipartite graphs. It is possible for the size of $\mathcal{P}^m_{\text{BIP}}(G)$ to be exponential in the size of G, even if G almost satisfies $\mathcal{P}_{\text{BIP}}$. Consider, for example, the graph G_2 in Figure 1. The graph G_2 almost satisfies $\mathcal{P}_{\text{BIP}}$ since $G_2 - w \in \mathcal{P}_{\text{BIP}}$. However, the set $\mathcal{P}^m(G_2)$ contains $2^n + 1$ graphs, i.e., $G_2 - w$ and the graphs derived by choosing the vertex w and one from each pair of vertices (v_i, u_i), for all i. Notwithstanding the size of $\mathcal{P}^m(G_2)$, it is not difficult to show that the input-restricted $\mathcal{P}_{\text{BIP}}$-subgraphs problem is P-enumerable. To see this, observe that it is possible to find the graphs in $\mathcal{P}^m(G_2)$ efficiently, in terms of the input *and* the output, by dealing separately with each connected component of $G_2 - w$. □

By formalizing the intuition presented in Example 4.1, the following propositions can be shown. Similar propositions can be shown for other graph properties.

Proposition 4.2. *The io-restricted $\mathcal{P}_{\text{CBIP}}$-subgraphs problem is in* PTIME.

Proposition 4.3. *The input-restricted $\mathcal{P}_{\text{BIP}}$-subgraphs is* P*-enumerable.*

5 Hereditary Properties

In this section, we reduce the maximal $\mathcal{P}$-subgraphs problem to the input-restricted maximal $\mathcal{P}$-subgraphs problem for hereditary properties $\mathcal{P}$. Our reduction is by means of an algorithm that shows how to compute $\mathcal{P}^m(G)$ for an arbitrary graph G, given a procedure that can compute $\mathcal{P}^m(G)$ for graphs G that almost satisfy $\mathcal{P}$.

In Figure 2, the algorithm GenHered$\langle\mathcal{P}\rangle$ is presented. This algorithm uses the following two procedures.

- Max$\langle\mathcal{P}\rangle(H, G)$: This procedure receives graphs H and G as input and returns true if H is a maximal $\mathcal{P}$-subgraph of G and false otherwise. This procedure can easily be defined in terms of Sat$\langle\mathcal{P}\rangle$, by *(1)* checking if $H \in \mathcal{P}$ *and (2)* extending H with each vertex in G and checking whether any extension is in $\mathcal{P}$ (using Sat$\langle\mathcal{P}\rangle$).

```
   Algorithm:    GenHered⟨P⟩
   Input:        Graph G = ({v_1, ..., v_n}, E, r)
   Output:       Maximal P-Subgraphs of G, i.e., P^m(G)
()
 1  G := {O_0}
 2  for i := 1 to n
 3  do H := G
 4       H ∈ HG := G − {H}
 5       H' ∈ GenRestrHered⟨P⟩(G[H, v_i])   if Max⟨P⟩(H', G[{v_1, ..., v_i}])
 6                                            then G := G ∪ {H'}
 7  return G
```

Figure 2: Algorithm to compute $\mathcal{P}^m(G)$ for hereditary properties

- GenRestrHered$\langle\mathcal{P}\rangle(H)$: This procedure receives a graph H that almost satisfies $\mathcal{P}$ and returns the set $\mathcal{P}^m(H)$. This procedure is not defined in this paper. Instead, it must be provided on a per-property basis.

In essence, our algorithm reduces the maximal $\mathcal{P}$-subgraphs problem to the input-restricted maximal $\mathcal{P}$-subgraphs problem by using the procedure GenRestrHered$\langle\mathcal{P}\rangle$.

The algorithm GenHered$\langle\mathcal{P}\rangle$ starts with the set $\mathcal{G} = \{O_0\}$, where O_0 is an empty graph, i.e., a graph with no vertices or edges. It then continuously (Line 6) attempts to extend each graph H in $\mathcal{G}$ with an additional vertex v_i. The graphs in $\mathcal{P}^m(G[H, v_i])$ that are maximal with respect to the vertices seen thus far are inserted into $\mathcal{G}$ (Line 7). This step is critical, since it *(1)* inserts graphs that are needed in order to create the final result *and (2)* avoids inserting extra graphs that would cause $\mathcal{G}$ to grow exponentially.

Let G be a graph with n vertices. Suppose that there are K graphs in $\mathcal{P}^m(G)$. We show that GenHered$\langle\mathcal{P}\rangle(G)$ correctly computes $\mathcal{P}^m(G)$ and analyze the runtime of our algorithm as a function of n and K. We use $s_{\mathcal{P}}(n)$ to denote the amount of time needed to check if $G \in \mathcal{P}$, i.e., the runtime of Sat$\langle\mathcal{P}\rangle(G)$. Observe that Max$\langle\mathcal{P}\rangle(H, G)$ runs in $\mathcal{O}(n\, s_{\mathcal{P}}(n))$ time. We use $r_{\mathcal{P}}(n, K)$ to denote the amount of time needed to compute $\mathcal{P}^m(G)$, when G almost satisfies $\mathcal{P}$, i.e., the runtime of GenRestrHered$\langle\mathcal{P}\rangle(G)$. Note that $r_{\mathcal{P}}$ is a function of both the input and the output.

Theorem 5.1. *Let $\mathcal{P}$ be a hereditary property and let G be a graph with n vertices. Let K be the number of graphs in $\mathcal{P}^m(G)$. Then*

- GenHered$\langle\mathcal{P}\rangle(G) = \mathcal{P}^m(G)$ and
- GenHered$\langle\mathcal{P}\rangle(G)$ *runs in time:* $\mathcal{O}\big(n^2\, s_{\mathcal{P}}(n)\, K\, r_{\mathcal{P}}(n, K)\big)$.

Proof (Sketch). We use G_i to denote the induced subgraph of G containing exactly the vertices $v_1, \ldots, v_i$, i.e., $G[\{v_1, \ldots, v_i\}]$. We use $\mathcal{G}_i$ to denote the value of $\mathcal{G}$ after i iterations of the loop in Line 2 of the algorithm. We show by induction on the number of vertices k in G that $\mathcal{G}_k = \mathcal{P}^m(G_k)$. The inclusion $\mathcal{G}_k \subseteq \mathcal{P}^m(G_k)$ can be shown by a case analysis of the lines in which graphs are added to (and removed from) $\mathcal{G}_k$. The inclusion $\mathcal{P}^m(G_k) \subseteq \mathcal{G}_k$ follows from the following inclusion:

$$\mathcal{P}^m(G_k) \subseteq \mathcal{P}^m(G_{k-1}) \cup \left(\cup_{H \in \mathcal{P}^m(G_{k-1})} \mathcal{P}^m(G[H, v_k]) \right).$$

The runtime follows from a careful analysis of the algorithm and from the fact that $|\mathcal{P}^m(G_{i-1})| \leq |\mathcal{P}^m(G_i)|$, for all i. □

Corollary 5.2. *Let $\mathcal{P}$ be a hereditary property. Then the maximal $\mathcal{P}$-subgraphs problem is in* PIO *if and only if the input-restricted maximal $\mathcal{P}$-subgraphs problem is in* PIO.

Corollary 5.3. *Let $\mathcal{P}$ be a hereditary property. Then the maximal $\mathcal{P}$-subgraphs problem is* P*-enumerable if the input-restricted maximal $\mathcal{P}$-subgraphs problem is in* PTIME.

Corollary 5.4. *The maximal $\mathcal{P}_{\text{BIP}}$-subgraphs problem is in* PIO.

6 Connected-Hereditary and Rooted-Hereditary Properties

GenHered$\langle\mathcal{P}\rangle$ may fail to return the correct graphs if $\mathcal{P}$ is connected hereditary or rooted hereditary. Intuitively, this failure is caused by the fact that an induced subgraph H may not be connected or rooted (and therefore, $H \notin \mathcal{P}$), even though there is a graph G such that $H \subset_{\text{IS}} G$ and $G \in \mathcal{P}$. In other words, the order in which we choose the vertices can effect the success of the algorithm.

In this section, we solve the maximal $\mathcal{P}$-subgraphs problem for connected-hereditary and rooted-hereditary properties in the following way:

- The maximal $\mathcal{P}$-subgraphs problem is reduced to the output-restricted maximal $\mathcal{P}$-subgraphs problem in a straightforward fashion.
- The output-restricted maximal $\mathcal{P}$-subgraphs problem is reduced to the io-restricted maximal $\mathcal{P}$-subgraphs problem by means of the algorithm GenWithVertex$\langle\mathcal{P}\rangle$ (see Figure 3). The result of calling GenWithVertex$\langle\mathcal{P}\rangle(G, v_r)$, for an arbitrary graph G, is the set $\mathcal{P}^m(G;\, v_r)$. Note that GenWithVertex$\langle\mathcal{P}\rangle(G, v_r)$ uses GenRestrWithVertex$\langle\mathcal{P}\rangle(G, v_r)$ which generates $\mathcal{P}^m(G';\, v_r)$ for graphs G' that almost satisfy $\mathcal{P}$.

In the remainder of this section, we explain the algorithm GenWithVertex$\langle\mathcal{P}\rangle$—its notation, data structures and flow of execution. We show its correctness and analyze its runtime.

Notation. Consider graphs G and H such that $H \subseteq_{\text{IS}} G$. We say that a vertex v in $V(G) - V(H)$ is an *undirected neighbor* of a vertex v' in H if either the edge (v, v') or the edge (v', v) is in G. Similarly, we say that $v \in V(G) - V(H)$ is a *directed neighbor* of v' in H if the edge (v', v) is in G. Note that the neighbors (directed or undirected) of vertices in an induced subgraph are *not* in the induced subgraph.

Given a property $\mathcal{P}$, we use $N_{\mathcal{P}}(H, G)$ to denote the set of undirected neighbors of H if $\mathcal{P}$ is connected-hereditary and to denote the set of directed neighbors of H if $\mathcal{P}$ is rooted-hereditary. Note that we use $\mathcal{P}$ only in order to differentiate between undirected and directed neighbors.

In our algorithm, a graph $H \subseteq_{\text{IS}} G$ (such that $H \in \mathcal{P}$) is associated with a set of vertices $\overline{V}(H)$. Intuitively, this set contains vertices $v' \in N_{\mathcal{P}}(H, G)$ that cannot be used to extend H, since $G[H, v'] \notin \mathcal{P}$.

Data Structures. GenWithVertex$\langle\mathcal{P}\rangle$ uses two stacks to collect graphs: $Stack_1$ and $Stack_2$. $Stack_1$ contains graphs for which processing is incomplete. Therefore, a graph H will be in $Stack_1$ if it has a $v' \in N_{\mathcal{P}}(H, G)$ that is not in $\overline{V}(H)$. For such a vertex v', it is not yet known whether v' can be added to H, i.e., whether $G[H, v] \in \mathcal{P}$. $Stack_2$ contains graphs for which processing is complete. Therefore, a graph H will be in $Stack_2$ if $N_{\mathcal{P}}(H, G) \subseteq \overline{V}(H)$.

To ensure that $Stack_1$ and $Stack_2$ contain the proper graphs, our algorithm uses the procedure PushAppropriate$\langle\mathcal{P}\rangle(H, G, Stack_1, Stack_2)$, which does the following. If $N_{\mathcal{P}}(H, G) \not\subseteq \overline{V}(H)$, then the procedure adds H to the top of $Stack_1$. Otherwise, the procedure adds H to the top of $Stack_2$.

Flow of Execution. The algorithm GenWithVertex$\langle\mathcal{P}\rangle$ starts by considering the graph $G[\{v_r\}]$. Then, it continually extends graphs in $Stack_1$ with neighboring vertices to derive larger graphs that are in $\mathcal{P}$. All extensions created must contain the vertex v_r (so that we will only create graphs in $\mathcal{P}^m(G;\, v_r)$). Suppose $H \in Stack_1$ and $v \in N_{\mathcal{P}}(H, G)$. We deal with the case in which $G[H, v] \in \mathcal{P}$ in Lines 9-10. We deal with the case in which $G[H, v] \notin \mathcal{P}$ in Lines 11-25.

Correctness and Runtime Analysis. The proof of correctness of GenWithVertex$\langle\mathcal{P}\rangle$ is rather intricate and has been omitted due to lack of space. However, we take note of the behavior of our algorithm that is critical in proving its correctness:

- The fact that $\overline{V}(H)$ is assigned the empty set every time that a vertex is added to H, allows us to be prove that *all* graphs in $\mathcal{P}^m(G;\, v_r)$ are returned.
- Graphs G' from $\mathcal{P}^m(G[H, v];\, v_r)$ (see Line 12) are not immediately added to $Stack_1$ or $Stack_2$. Instead we first try to combine G' with existing graphs in $Stack_1$. We also check if G' is an induced subgraph of a graph in $Stack_2$. Only if we have not succeeded in either of these actions, do we add G' to $Stack_1$ or $Stack_2$, as appropriate. This prevents $Stack_1$ and $Stack_2$ from growing too big. It also ensures that *only* graphs from $\mathcal{P}^m(G;\, v_r)$ are returned, and each such graph is returned only once.

In order to prove our complexity analysis of the runtime of GenWithVertex$\langle\mathcal{P}\rangle$, we must show that GenRestrWithVertex$\langle\mathcal{P}\rangle$ does not create more graphs than the number of graphs in the result of GenWithVertex$\langle\mathcal{P}\rangle$. This holds because we are able to prove that $H \subset_{\text{IS}} G$ implies that $|\mathcal{P}^m(G[H];\, v_r)| \leq |\mathcal{P}^m(G;\, v_r)|$ if either *(1)* $\mathcal{P}$ is connected-hereditary *or (2)* $\mathcal{P}$ is rooted-hereditary and v_r is the root of G.

Let G be a graph with n vertices. We use $r'_{\mathcal{P}}(n, K)$ to denote the amount of time needed to compute $\mathcal{P}^m(G;\, v)$ for a graph G that almost satisfies $\mathcal{P}$ and an arbitrary vertex v, i.e., the runtime of the procedure GenRestrWithVertex$\langle\mathcal{P}\rangle(G, v)$. Note that $r'_{\mathcal{P}}$ is a function of both the input n and the number of graphs in the output K. The function $s_{\mathcal{P}}$ is defined as before.

Theorem 6.1. *Let $\mathcal{P}$ be a connected-hereditary or rooted-hereditary property. Let G be a graph with n vertices, and let v_r be a vertex in G. Suppose that $G[\{v_r\}] \in \mathcal{P}$. Let K be the number of graphs in $\mathcal{P}^m(G;\, v)$. Then*

- GenWithVertex$\langle\mathcal{P}\rangle(G, v_r) = \mathcal{P}^m(G;\, v_r)$ and
- GenWithVertex$\langle\mathcal{P}\rangle(G, v_r)$ *runs in time:* $\mathcal{O}\Big(n^2\, K^2\, r'_{\mathcal{P}}(n, K)\big(s_{\mathcal{P}}(n) + n\big)\Big)$.

```
Algorithm:   GenWithVertex⟨𝒫⟩
Input:       Graph G and Vertex v_r
Output:      Maximal answers that contain vertex v_r, i.e., 𝒫^m(G; v_r)
()
 1  V̄(G[{v_r}]) := ∅
 2  Stack_1 := ∅
 3  Stack_2 := ∅
 4  PushAppropriate⟨𝒫⟩(G[{v_r}], G, Stack_1, Stack_2)
 5  while Stack_1 ≠ ∅
 6  do H := Stack_1.Pop()
 7     let v be a vertex in N_𝒫(H, G) − V̄(H)
 8     if Sat⟨𝒫⟩(G[H, v])
 9       then V̄(G[H, v]) := ∅
10            PushAppropriate⟨𝒫⟩(G[H, v])
11       else V̄(H) := V̄(H) ∪ {v}
12            G' ∈ GenRestrWithVertex⟨𝒫⟩(G[H, v], v_r) − {H}V̄(G') := ∅
13            inserted := false
14            H' ∈ Stack_1 s.t. Sat⟨𝒫⟩(G[G', H'])G_new := G[G', H']
15            V̄(G_new) := ∅
16            Stack_1.Remove(H')
17            PushAppropriate⟨𝒫⟩(G_new, G, Stack_1, Stack_2)
18            inserted := true
19            if exists H' ∈ Stack_2 s.t. V(G') ⊆ V(H')
20              then inserted := true
21            if not(inserted)
22              then PushAppropriate⟨𝒫⟩(G', G, Stack_1, Stack_2)
23            PushAppropriate⟨𝒫⟩(H, G, Stack_1, Stack_2)
24  return Stack_2
```

Figure 3: Algorithm to compute $\mathcal{P}^m(G;\ v_r)$ for connected-hereditary or rooted-hereditary properties

Theorem 6.2. *Let $\mathcal{P}$ be a connected-hereditary or rooted-hereditary property. The maximal $\mathcal{P}$-subgraphs problem is in* PIO *if the io-restricted maximal $\mathcal{P}$-subgraphs problem is in* PIO*.*

Proof (Sketch). Since any vertex in G may or may not appear in a solution of a connected-hereditary property, it is possible to compute $\mathcal{P}^m(G)$ for connected-hereditary properties by calling GenWithVertex$\langle\mathcal{P}\rangle$ for every vertex in G. If $\mathcal{P}$ is a rooted-hereditary, every graph in $\mathcal{P}^m(G)$ must contain the root of G. Hence, it is possible to compute $\mathcal{P}^m(G)$ by calling GenWithVertex$\langle\mathcal{P}\rangle$ with the root of G. □

If $\mathcal{P}$ is rooted hereditary and v_r is the root of G, then $\mathcal{P}^m(G;\ v_r) = \mathcal{P}^m(G)$. Hence, we can show the following result.

Corollary 6.3. *Let $\mathcal{P}$ be a rooted-hereditary property. The maximal $\mathcal{P}$-subgraphs problem is in* PIO *if and only if the input-restricted maximal $\mathcal{P}$-subgraphs problem is in* PIO*.*

7 Extending the Algorithms

In this section we discuss some small changes that can be made to the algorithms GenHered$\langle\mathcal{P}\rangle$ and GenWithVertex$\langle\mathcal{P}\rangle$ in order to improve the complexity results from the previous sections.

P-Enumerable for Connected-Hereditary and Rooted-Hereditary Properties. In Corollary 5.3, we presented a sufficient condition for the maximal $\mathcal{P}$-subgraphs problem to be P-enumerable, for hereditary properties $\mathcal{P}$. The algorithm GenWithVertex$\langle\mathcal{P}\rangle$ cannot be used in order to derive a sufficient condition for this problem to be P-enumerable for connected-hereditary or rooted-hereditary properties, since K appears quadratically in the runtime of GenWithVertex$\langle\mathcal{P}\rangle$.

It turns out that one can adapt GenHered$\langle\mathcal{P}\rangle$ to derive an algorithm that computes $\mathcal{P}^m(G;\, v_r)$, for a rooted-hereditary or connected-hereditary property $\mathcal{P}$, provided that certain conditions hold. The crux of the change to GenHered$\langle\mathcal{P}\rangle$ is in careful choice of the order in which to iterate over the vertices in G. The adapted algorithm can be used similarly to GenWithVertex$\langle\mathcal{P}\rangle$ in order to compute $\mathcal{P}^m(G)$.

Theorem 7.1. *Suppose that the input-restricted maximal $\mathcal{P}$-subgraphs problem is in* PTIME. *Then, the maximal $\mathcal{P}$-subgraphs problem is* P*-enumerable if* (1) *$\mathcal{P}$ is rooted-hereditary and G is acyclic* or (2) *$\mathcal{P}$ is connected-hereditary and the underlying undirected graph of G is a tree.*

Incremental Polynomial Time. None of the complexity results presented have provided conditions for the maximal $\mathcal{P}$-subgraphs problem to be solvable in incremental polynomial time. By slightly changing the procedure GenWithVertex$\langle\mathcal{P}\rangle$ we can derive an algorithm that computes $\mathcal{P}^m(G;\, v_r)$ in incremental polynomial time for an important special case. Using this adapted algorithm, $\mathcal{P}^m(G)$ can also be computed in incremental polynomial time, for connected-hereditary and rooted-hereditary properties. Our adapted algorithm can also be used for a hereditary property $\mathcal{P}$, by reducing $\mathcal{P}$ to an appropriately defined rooted-hereditary property. In addition, we derive a polynomial complexity result for returning k maximal induced subgraphs, for any constant k.

Theorem 7.2. *Let $\mathcal{P}$ be hereditary, connected-hereditary or rooted-hereditary property. Suppose that the io-restricted maximal $\mathcal{P}$-subgraphs problem is in* PTIME. *Then,* (1) *the maximal $\mathcal{P}$-subgraphs problem is in* PINC and (2) *k graphs from $\mathcal{P}^m(G)$ can be returned in polynomial time, for any constant k.*

Proof (Sketch). Item 2 follows directly from Item 1. To show Item 1, let G be a graph with n vertices. By careful observation, one may note that after at most n^2 iterations of the loop in Line 5 of GenWithVertex$\langle\mathcal{P}\rangle$, an additional graph will be in $Stack_2$. One can take advantage of this fact to adapt GenWithVertex$\langle\mathcal{P}\rangle$ so that it will run in incremental polynomial time, by having PushAppropriate$\langle\mathcal{P}\rangle$ print graphs as it adds them to $Stack_2$. (Care has to be taken not to print graphs that appeared before in a previous call to GenWithVertex$\langle\mathcal{P}\rangle$.) □

Corollary 7.3. *The maximal $\mathcal{P}_{\text{CBIP}}$ problem is in* PINC.

8 Conclusion

This paper reduces the maximal $\mathcal{P}$-subgraphs problem to restricted versions of the same problem by providing algorithms that solve the general problem, assuming that an algorithm for a restricted

version is given. Our results imply that when attempting to efficiently solve the maximal $\mathcal{P}$-subgraphs problem, it is not necessary to define an algorithm that works for the general case. Instead, an algorithm for restricted cases must be defined. An efficient method for solving the maximal $\mathcal{P}$-subgraphs problem for the general case is automatically derived from our algorithms.

Sometimes it turns out that algorithms for restricted cases of the maximal $\mathcal{P}$-subgraphs problem are straightfoward. For example, this is the case with the properties $\mathcal{P}_{\text{BIP}}$ and $\mathcal{P}_{\text{CBIP}}$. There are additional properties for which this holds, e.g., the set of independent sets, the set of star graphs, etc. Thus, our results immediately imply that the maximal "is an independent set"-subgraphs problem is both P-enumerable and in PINC, and the maximal "is a star graph"-subgraphs problem is in PINC. Note that it is significantly easier to come up with algorithms that solve the restricted versions of these problems than to come up with algorithms that solve the general cases.

Interestingly, our results can be applied to the database problem of computing maximal query answers. Well-known semantics for this problem, e.g., full disjunctions [7], can be modeled as graph properties. It is often easy to define algorithms that solve the restricted versions of the maximal $\mathcal{P}$-subgraph problem, for graph properties that correspond to semantics for incomplete information. Hence, the results in this paper have immediate practical applications for efficiently computing maximal query answers.

References

[1] E. Akkoyunlu. The enumeration of maximal cliques of large graphs. *SIAM Journal on Computing*, 2:1–6, 1973.

[2] D. Babic and A. Graovac. Enumeration of acyclic walks in a graph. *Discrete Applied Mathematics*, 45:117–123, 1993.

[3] J. Balogh, B. Bollobás, and D. Weinreich. The speed of growth for hereditary graph properties. *Journal of Combinatorial Theory*, 79:131–156, 2000.

[4] S. Cohen, Y. Kanza, and Y. Sagiv. Generating relations from XML documents. In *Proc 9th International Conference on Database Theory*, Siena (Italy), Jan. 2003. Springer-Verlag.

[5] S. Cohen and Y. Sagiv. An abstract framework for generating all maximal answers to queries, 2004. Manuscript available at `http://www.cs.huji.ac.il/~sarina/db_hereditary.ps`.

[6] K. Fukuda, T. Matsui, and Y. Matsui. Algorithms for combinatorial enumeration problems, 1995. `http://dmawww.epfl.ch/roso.mosaic/kf/enum/comb/combenum`.

[7] C. Galindo-Legaria. Outerjoins as disjunctions. In *Proc. 1994 ACM SIGMOD International Conference on Management of Data*, pages 348–358, Minneapolis (Minnesota, USA), May 1994. ACM Press.

[8] M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, San Francisco, 1979.

[9] D. Johnson. Finding all the elementary circuits of a directed graph. *SIAM Journal on Computing*, 4:77–84, 1975.

[10] D. Johnson and M. Yannakakis. On generating all maximal independent sets. *Information Processing Letters*, 27(3):119–123, 1988.

[11] Y. Kanza, W. Nutt, and Y. Sagiv. Queries with incomplete answers over semistructured data. In *Proc. 18th Symposium on Principles of Database Systems*, Philadelphia (Pennsylvania, USA), May 1999. ACM Press.

[12] Y. Kanza and Y. Sagiv. Computing full disjunctions. In *Proc. 22nd Symposium on Principles of Database Systems*, pages 78–89, San Diego (California, USA), June 2003. ACM Press.

[13] S. Kapoor and H. Ramesh. Algorithms for enumerating all spanning trees of undirected and weighted graphs. *SIAM Journal on Computing*, 24(2):247–265, 1995.

[14] J. Lewis and M. Yannakakis. The node-deletion problem for hereditary properties is NP-complete. *Journal of Computer and System Sciences*, 20(2):219–230, 1980.

[15] B. Madsin, J. Nielsen, and B. Skjernaa. On the number of maximal independent sets in a graph. Technical Report RS-02-17, Aarhus Univ., Datalogisk afdeling, Center for Basic Research in Computer Science (BRICS), April 2002.

[16] R. Read and R. Tarjan. Bounds on backtrack algorithms for listing cycles, paths, and spanning trees. *Networks*, 5:237–252, 1975.

[17] V. Sos and M. Saks. The diversity of hereditery graph properties. In *Combinatorics in Oporto*, Porto (Portugal), Sept. 2003.

[18] S. Tsukiyama, M. Ide, H. Ariyoshi, and I. Shirakawa. A new algorithm for generating all maximal independant sets. *SIAM Journal on Computing*, 6:505–517, 1977.

[19] L. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8:189–201, 1979.

[20] M. Yannakakis. Node- and edge-deletion NP-complete problems. In *Proc. 10th Annual ACM Symposium on Theory of Computing*, pages 253–264, New York (USA), 1978. ACM Press.

[21] M. Yannakakis. Algorithms for acyclic database schemas. In W. Chu, G. Gardarin, S. Ohsuga, and Y. Kambayashi, editors, *Proc. 7th International Conference on Very Large Data Bases*, pages 82–94, Cannes (France), Sept. 1981. Morgan Kaufmann Publishers.